

## ISDOCID: EW0 9821660A1 | &gt;

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

## DYNAMICALLY UPGRADEABLE DISK ARRAY SYSTEM AND METHOD

Technical Field

This invention relates to dynamically adding a disk array chassis to an already  
5 operating disk drive array, and communicating with the new disk array chassis over a  
communication path before the disk drives are interfaced on a high speed path with the  
disk drive controller.

Background of the Invention

Subsystems comprising disk arrays, i.e., groups of small, independent disk drive  
10 modules used to store large quantities of data have been developed and found to possess  
many advantages over a single large disk drive. For example, the individual modules of a  
disk array typically take up very little space and typically use less power and cost less than  
a single large disk drive, yet, when grouped together in an array, provide the same data  
storage capacity as a single large disk drive. In addition, the small disks of an array  
15 retrieve data more quickly than does a single large disk drive because, with a small disk  
drive, there is less distance for the actuator to travel and less data per individual disk to  
search through. The greatest advantage to small disk drives, however, is the boost they  
give to I/O performance when configured as a disk array subsystem.

A disk array is typically housed in a chassis, which may be an enclosure, for  
20 holding a plurality of disk drives. When additional memory storage space is required, it is  
desirable to be able to add an additional disk array chassis to the existing system. It is  
further desirable that one be able to add the new disk array chassis to the existing system  
without needing to shut down the existing system. In other words, it is desirable to  
provide a dynamically upgradeable disk array system.

25 For a system including a disk controller communicating over a loop with a  
plurality of disks, it is undesirable to expand that bus by simply adding more disk drives.  
There may be a malfunction in any of the new disk drives. Further there may be a  
problem with the new connection to the existing loop. While the functional status of the  
new equipment can be tested by the disk controller after it is connected to the loop, if  
30 there is a malfunction, the entire loop (and disk drives) will be degraded and may be

-2-

rendered inoperable or inaccessible. Therefore, it is desirable to be able to determine the integrity of the new disk drives and bus before permitting interface with the existing system.

### Summary of the Invention

5 In accordance with an embodiment of the present invention, a dynamically upgradeable disk array chassis includes a chassis that contains the disk drives. The disk drives and a disk controller are connected in a loop with a first serial bus for transmitting data from the controller and a second serial bus for passing data back to the controller. The disk drives are connected to one or the other of the serial busses. A shunt is  
10 connected to the first and second busses at the end opposite the controller and has a normal state in which the first bus is connected to the second bus thus completing the loop. The shunt may also assume a state in which each of the first and second serial busses are connected to separate outputs. The separate outputs are used to connect to a new disk array chassis with similar busses and shunting device. Each chassis has an  
15 environmental monitor that controls the state of its shunt. A separate serial communication bus connects to the monitor. When a new disk array chassis is connected to the existing chassis, the existing chassis's environmental monitor may communicate through its serial communication bus with the serial communication bus and environmental monitor on the new disk array chassis to determine whether conditions are  
20 suitable for extending the disk drive loop to include the disks on the new disk array chassis. If no faults are detected, the existing chassis's environmental monitor switches its shunt to connect the disk drive serial busses of the existing disk array chassis to those of the new disk array chassis. The new chassis's environmental monitor leaves its shunt in the normal state to connect the two disk drive busses and the disk drive loop is again  
25 complete.

In accordance with an embodiment of the invention, diplexing is used to provide the path for the serial communication bus to the environmental monitor, and the signals for the disk drive loop over the same wires of a single cable connecting the disk array chassis to one another. Preferably, low frequency signals are used for the communication  
30 with the environmental monitor and high frequency signals are used for disk drive loop communications including I/O operations. The low frequency signals may be similar to

those of the RS-232 serial data protocol and the high frequency signals may adopt the Fibre Channel Arbitrated Loop signal protocol. The Fibre Channel 8B/10B signal encoding advantageously guarantees frequent transitions on the Fibre Channel signal so that the signal can always be distinguished from the low frequency signals. The Fibre Channel encoding keeps the frequency content of the Fibre Channel signal isolated to 100 MHZ and above by guaranteeing a transition at least every 5 bits.

A diplexer is connected to the environmental monitor serial communication bus and the separate disk drive serial bus outputs of the shunt for combining the high frequency data signals from the disk drive controller and the low frequency data signals from the environmental monitor serial communication bus which are to be sent out from the chassis. The diplexer also includes a configuration for separating the high frequency disk drive serial bus data signals from the low frequency environmental monitor SC8 data signals coming in from the new disk array chassis. By providing a single connector and diplexing and shunt control in the presently preferred embodiment for both the low and high frequency communication signals, the integrity of the new chassis may be determined over the low frequency environmental monitor communication path without electrically connecting the high frequency path and possibly disrupting operation on the existing disk array system. Only after the integrity of the connection and the operability of the new disk array enclosure is determined will the existing chassis shunt be switched into the state permitting connection of the high frequency data path with the disk drives of the new disk array chassis.

Referring now more particularly to the multiplexing of the two signals, it is noted that a common Fibre Channel (FC) data communications link may be used. In preferred embodiments of the invention, the first signal to be diplexed is a low frequency asynchronous or synchronous RS-232 type serial signal; and the second signal is the standard high frequency Fibre Channel signal. The terms "low frequency" and "high frequency" are used herein as relative to one another. As used herein, a low frequency signal is one that can be separated by filtering from the high frequency signal and vice versa. The low frequency signal, which in preferred embodiments is used to communicate between environmental monitor UARTS, is extremely slow in comparison to the Fibre Channel signal. This signal, preferably transmitted at 9600 baud, has a maximum

fundamental frequency of 4.8 kHz, and a minimum frequency of DC. The square wave nature of this signal will normally have higher order harmonics which extend to very high frequencies. Consequently, in preferred embodiments of the invention, these harmonics may be limited by low pass filtering without affecting the information content of the signal, with the intended result of isolating the RS-232 type signal to a low frequency band from DC to about 100 kHz. The high frequency FC signal is limited to 100 MHz, and beyond by the aforementioned 8B/10B encoding.

In an embodiment of the present invention, the high frequency Fibre Channel, and the low frequency RS-232 type signal, are frequency division multiplexed into a composite signal which is sent over a single cable. The cable will preferably be a twinax Fibre Channel arbitrated loop standard configuration having two conductor pairs (1 pair for each direction) contained within a common shield. Isolation between the two conductor pairs is maintained by locating the two pairs orthogonally in the cable. Each pair will contain a FC signal in differential mode, and a RS-232 signal in common mode. In each conductor pair, one wire conducts a sum of the high frequency FC signal and the low frequency RS-232 type signal and the other wire conducts the difference resulting when the high frequency FC signal is subtracted from the low frequency RS-232 type signal.

20

#### Brief Description of the Drawings

The invention may be more readily understood by reference to the following detailed description, taken with the accompanying drawings, in which:

FIG. 1 is a block diagram of a chain of dynamically upgradeable disk array chassis of the present invention.

25

FIG. 2 is a flow chart of an upgrading algorithm for use with the present invention.

FIG. 3 is a block diagram of a system of interconnected disk array chassis.

FIG. 4 is a block diagram of an environmental monitor for use in the disk array chassis of FIG. 1.

FIG. 5 is a graph showing the duplex filter frequency response.

30

FIG. 6 is a figure showing the composite Fibre Channel and RS-232 type signals.

FIG. 7 is a circuit diagram for the differential coupling circuit.

FIG. 8 is a circuit diagram showing the common mode coupling circuit.

FIG. 9 is a functional diagram showing an embodiment of the invention using twin-ax or twisted-pair cable and common mode coupling.

FIG. 10 shows an implementation of the FIG. 9 circuit diagram implemented using  
5 transformers.

FIG. 11 is a schematic circuit diagram of the low frequency signal input and output portions of a diplexer for use in the disk array chassis of FIG. 1.

FIG. 12 is a schematic circuit diagram of the high frequency signal input portion of a diplexer for use in the disk array chassis of FIG. 1.

10 FIG. 13 is a schematic circuit diagram of the high frequency signal output portion, a shunt and an expansion connector for use in the disk array chassis of FIG. 1.

#### Detailed Description of the Preferred Embodiments

Referring now to the drawings, a disk array chassis 10 is shown in block diagram  
15 form in FIG. 1. The disk array chassis provides a series of openings into which disk drives 12 may be inserted. The chassis 10 of FIG. 1 is shown capable of housing a specified number of disk drives 12, for example ten. The mechanical chassis may be an open or closed housing or enclosure. Inside the chassis are conventional auxiliary equipment 14 including cooling fans, power supplies and a display card. The display  
20 card controls a display on the outside of the housing which may be used to exhibit certain status conditions. Status conditions are determined by the environmental monitor.

The disk drives are interconnected by a high speed serial data communication path. In accordance with the presently preferred embodiment, the communication path is a loop including a first serial bus 22 directing data in one direction and a second serial bus 24  
25 directing data in a second opposite direction. The preferred data loop is a Fibre Channel arbitrated loop. A series of port bypass circuits 26 are inserted along the loop for making the connection with the disk drives 12. Each port bypass circuit 26 is controlled by an environmental monitor 30. The control line between the environmental monitor 30 and the port bypass circuit 26 exists, but is not shown in FIG. 1. When the port bypass circuit  
30 is in its bypassed condition, its associated disk drive is shunted out of the loop. The environmental monitor 30 switches the port bypass circuit 26 into a non-bypassed

condition so that signals on the serial bus are directed to and from the associated disk drive. In order to provide a high availability system, the disk array chassis of the presently preferred embodiment may include two high speed data communication loops and environmental monitors acting redundantly. Each loop is connected ultimately to a disk controller 16. The two disk control processors 16, one for each loop, further serve to provide high availability of the data storage system.

It is desirable that the disk array chassis permit dynamic upgrading. In that regard, a shunt 40 is provided along the Fibre Channel loop. In the presently preferred embodiment, the shunt 40 may be a port bypass circuit. The first serial bus 22 providing data in a first direction is connected to the shunt and the second serial bus 24 providing data in the opposite direction is also connected to the shunt. When the shunt is in its normally bypassed condition, the Fibre Channel loop ends at the shunt and data passes directly from the first serial bus 22 through the shunt to the second serial bus 24. When the shunt is switched into a non-bypassed state, the first bus 22 is allowed to pass to an expansion connection 52 and the second bus is connected to an incoming signal from the expansion connection. This extends the Fibre Channel loop into the next chassis.

The shunt is controlled by the environmental monitor 30. The environmental monitor 30 is provided with a serial communication path 42 that can lead to an adjacent disk array chassis. The environmental monitor is a supplier of status signals and other environmental information. The monitor may be any entity that sequences the status signals onto the serial communication path 42. The environmental monitors of the presently preferred embodiment are provided with two universal asynchronous receiver transmitters 44 (UART) for communicating with a chassis connected at either end of the chassis housing the monitor. While the communication path 42 may be provided for connection to an adjoining disk array chassis separate from the Fibre Channel loop, in accordance with the presently preferred embodiment of the invention, the communication path 42 is provided as a path for low frequency signals which are diplexed with the high frequency Fibre Channel signals for transmission over a single cable 18. Preferably a twinax cable 18 is used. This is a standard cable for Fibre Channel. The present invention advantageously adds a low frequency signal to the high frequency Fibre Channel signals normally carried by the cable. A twinax cable 18 includes four wire



conductors, each divided into two pairs. Each pair transports signals in a single direction so one pair is used for the exiting signals and the other pair of wires in the twinax cable 18 is used for the returning signals.

A diplexer 50 is preferably connected to transmit and receive the low frequency signals from the environmental monitor 30 and also to transmit and receive high frequency signals from the shunt 40. The other end of the diplexer is connected to a cable connector 52. Regardless of the state of the shunt, communications between disk array chassis environmental monitors are possible over the single twinax cable 18. The disk array chassis includes a diplexer on both ends of the high frequency drive serial busses, except for the first chassis 110 that houses the disk controller 16 where only one diplexer is required. Communications can thus be maintained with the preceding and subsequent disk array chassis. The environmental monitor communicates with low frequency signals that can be separated from the high frequency Fibre Channel signals by filtering in the diplexer. The signals from the diplexed cables can be split with one copy of the signals being high pass filtered while the other copy of the signals is low pass filtered to extract each of the two separate signals. A preferred method of diplexing, however, will be described later herein.

A data storage system is configured by providing 1 or more disk control processors 16, each in communication with one of the Fibre Channel loops and in low frequency communication with the environmental monitors. Disk array chassis are connected to the processors in a continuous chain. The arrangement of the high frequency communication path with two serial busses through each disk array chassis providing communications in each direction permits a high frequency loop. The shunt in each disk array chassis is in the non-bypassed position except for the final active disk array chassis in the chain which remains in the normal bypass condition closing the loop. The disk control processors 16 may be provided in a first chassis 110 along with a disk array. This first chassis 110 in the chain only requires a single diplexer since there is no connection to a preceding chassis.

Dynamically upgrading a data storage system of the invention shall now be described. An existing chain of disk storage chassis is provided. The system may include a single chassis enclosing a disk array with the disk controller. Alternatively, the system

may include additional disk array chassis connected in a chain to the controller chassis. The system is upgraded by connecting an additional disk array chassis to the last disk array chassis in the existing chain. A communication cable 18 is connected from the new disk array chassis to the expansion connector 52 on the chassis of the last disk array chassis in the existing chain. This last disk array chassis in the chain has a shunt 40 which is in the normally bypassed condition connecting its first serial bus 22 to the second serial bus 24 and completing the Fibre Channel loop. When the cable 18 is inserted into the connector, the shunt remains in the bypassed condition. The new disk array chassis is now in communication via the low frequency signals conducted over the low frequency serial communication path with the environmental monitor 30. After the connection has been made, the environmental monitor 30 in the existing system can communicate with the environmental monitor 30 in the new disk array chassis. If it is determined that the new disk array chassis is properly connected and operating properly, the shunt 40 is switched into the non-bypassed state extending the high frequency communications loop onto the new chassis.

Referring now to FIG. 2, a flow chart of the communications conducted along the low frequency environmental monitor serial communication path, shall now be described. In accordance with the presently preferred embodiment, the commands supported by the serial protocol along the communication path include a poll, read and write command. These commands are issued by the disk controller 16. The poll command returns the address of each chassis and an indication if that enclosure has changed state since the last poll was issued. The address of a chassis is known to the environmental monitor 30. At present, it is preferred that the address be set by manual switches on the outside of the chassis. In this manner, a chassis can be given a distinct address different from each of the other chassis that have already been connected to the existing system. Preferably, the user will address the chassis in a numerical order such as 1, 2, 3, 4.... A read command is directed to a specific chassis address and returns detailed environmental status about the chassis including its power, cooling and disk drive status. A write command is used by the disk controller to control the operation of the chassis. This command allows the disk controller to turn on or off any port bypass circuit and the shunt 40 of any chassis.

With the shunt of the penultimate chassis still in the bypass condition, the disk

controller is able to monitor and control the newly added last chassis via the diplexed serial communication path 42 without disrupting the ongoing activity on the Fibre Channel loop. This allows the disk controller to verify the configuration and environmental status of the new chassis before enabling the new drives onto the Fibre Channel loop. As set forth in FIG. 2, the disk controller runs a periodic polling 60 of all the environmental monitors on the system over the low frequency serial communication path 42. The poll command goes out from the disk controller and is received by all of the environmental monitors that are connected to the low frequency serial communication path. Each environmental monitor 30 on the path returns its address and an indication of whether there has been any state change since the last poll. If a new chassis is detected 62 during this periodic polling process, a dynamic upgrading algorithm is initiated. A time delay 64 is imposed to allow for the configuration of the system to settle as the new chassis is powered up.

After completion of the time delay, the disk controller reads the status 68 from every environmental monitor 30 along the low frequency serial communication path 42. The read command produces from each environmental monitor information about the cabling and address of the chassis 70. This is used to verify that the cabling is properly configured. The address is checked to be sure it is within the range of legally assignable addresses. The read will also return information as to whether the current chassis is operational 72, i.e., powered and functioning properly. If a problem is found with respect to any of the environmental monitors, the Fibre Channel loop will be terminated at the previous chassis 74, 78 so that only operational chassis are permitted on the Fibre Channel loop. If the chassis is properly connected and operational, processing continues 76 until the new chassis is found 66. If the new enclosure is properly connected and is operational, then it will be set as the last enclosure in the system 76. The shunt 40 of the previous chassis will be switched into a switched state 80 to permit Fibre Channel communications with the disks on the new chassis. The Fibre Channel loop may then be instructed to either automatically if it is so set or upon manual indication perform a discovery process to reinitialize the operation on the Fibre Channel loop to include the new disks. This discovery process on the Fibre Channel loop is disruptive to any Fibre Channel activity and may take several seconds to complete. Therefore, it is preferred that

the system provide the system administrator with the option of automatically entering discovery or selecting a manual process that is initiated by a directive from the host system. The discovery process makes the disk drives known to the disk controller's Fibre Channel driver and allows the upper layer protocol drivers (SCSI) to communicate with  
5 the disks.

Advantageously, communications with the new disk array chassis along the high frequency Fibre Channel do not begin until the disk controller has had an opportunity to communicate with the new disk array chassis through the low frequency communication path. This process operates separately and independently through each of the two cables  
10 in a redundant system such as that shown in FIG. 3. Thus, a new disk array chassis in the redundant system embodiment is added by connecting two cables to the existing chain, one for each loop. The connection of each cable operates separately to enable communications over the respective Fibre Channel loop. FIG. 3 illustrates a chain of disk array chassis. At the bottom of the chain, the chassis 110 houses a pair of disk controllers  
15 as well as a disk array. Each disk controller is in connection with one of the Fibre Channel loops. The chassis with the disk controller is used for connection to one or more host bus adaptors 112. A host bus adaptor links a host computer with the memory storage provided by the disk arrays. Each of the disk array chassis 10 in the chain is shown with one of two cables connecting each of its redundant loops to the adjoining disk array  
20 chassis.

A presently preferred environmental monitor 30 for use in the invention is shown in greater detail in the block diagram of FIG. 4. The monitor program is run on a microcomputer 90. The microcomputer is shown connected to the dual UART's 44. One  
25 UART is connected to the input end of the chassis, the other is connected to the expansion end. The UART receives and transmits signals on the low frequency serial communication path 42.

The environmental monitor 30 also includes a number of status ports 100 for monitoring various conditions within the enclosure. One of the status signals is the chassis address (Encl Addr). According to the present embodiment, the address of the  
30 chassis is manually set by switches on the outside of the chassis. Thus, a chassis can be manually set to have an address different from the already present disk array chassis in the

existing chain of chassis. Other status signals include power system status, cooling status and drive status. The environmental monitor also includes port bypass circuit control latches 102. These control the port bypass circuits 26 for each of the disk drives that may be inserted into the disk array enclosure. Also one of the latches will control the shunt 40.

5 In preferred embodiments of the invention, low frequency signals and high frequency signals are frequency division multiplexed into a composite signal for transmission on cables between chassis. They may be simply combined and separated by a diplex filter at each end of the cable, hence use of the term "diplexing." FIG. 5 shows the diplex filter frequency response. FIG. 6 shows a RS-232 type signal, a FC signal, and  
10 the diplexed composite of the two. The diplex filter has three ports, the cable port, a low pass port and a high pass port. The filter is constructed by connecting the output of a high pass network and a low pass network together. The high pass network will preferably act as an open circuit at low frequency; the simplest form of this high pass network is a capacitor. The low pass network will preferably look like an open circuit to the high  
15 frequency Fibre Channel signal; the simplest form of the low pass network is an inductor or a resistor. The resistor may be used if the system can tolerate the DC resistance; alternatively, inductors may be used without loss of effectiveness, but they are not ideal devices since achieving broad bandwidth isolation with an inductor is more difficult than with resistors.

20 Diplexing may be used for communication between a variety of types of computer components. The connection between these components will preferably be with a pair of twinax cables carrying the two Fibre Channel loops. There is also an RS-232 type communications channel between environmental monitors. In embodiments of the invention, the use of twinax affords two options for coupling the RS-232 type low  
25 frequency signal onto the cables. In one embodiment of the invention, the signal may be applied as a differential signal on the twinax cable; in a preferred embodiment of the invention, the signal may be applied as a common mode signal to both sides of the twinax.

A circuit for the first method, differential coupling, is shown in FIG. 7. Coupling  
30 the RS-232 type low frequency signal differentially onto the twinax is one possible way to pass the signal from chassis to chassis. This is the same coupling method as that used for

the Fibre Channel signals. This method has the highest level of noise immunity for the low frequency signals. Noise sources such as ground noise and low frequency magnetic fields can be easily rejected using this technique. On the down side, this coupling method will not work if equalized cable assemblies are used (that's not to say that equalization cannot be done). Another detractor is that any residual coupling of the RS-232 type low frequency signal into the FC signal will be a differential noise source and cause some amount of jitter.

Regarding ground noise, the chassis to chassis twinax-interconnects will preferably be in the range of 0 to 10 meters in length. These connections will be between chassis within a rack, or between racks in close proximity. There may be a ground voltage potential between two systems in different racks. This ground noise is superimposed on the cable as a common mode signal. The differential receiver will reject the common mode noise, and properly receive the differential signal.

Regarding magnetic susceptibility, in an embodiment of the invention, the twinax cable will preferably be shielded and, in the preferred embodiment of the invention, the shield is connected to chassis ground at both ends. Although this configuration is not the optimal configuration for magnetic shielding, this is required for good high frequency emissions performance. If the cable is brought into an AC magnetic field, a current will be induced to flow through the cable shield, around the ground loop. This current will couple a noise voltage into the internal signals through the mutual inductance between the shield and signal conductors. This noise voltage will be coupled in as a common mode noise source in this configuration, and be rejected by the receiver.

Regarding coupling noise, in an embodiment of the invention, the Fibre Channel signal is very low amplitude relative to the RS-232 type low frequency signal. With the use of a simple R/C pulse shaping filter, and simple capacitive/resistive duplex filter, there is some residual RS-232 signal coupled into the Fibre Channel receiver as shown in FIG. 5. This noise is minimized by making the slew rate of the RS-232 type low frequency signal as slow as possible for the required data rate. The residual noise coupled into the FC receiver is differential noise in this instance, and results in some additive jitter on the received FC signal. This noise is manageable for data rates of 9600 baud and below, while better lowpass filtering may be required for higher data rates.

Regarding external compatibility, in an embodiment of the invention, the rejection of the low frequency signal in the Fibre Channel receiver is accomplished by the AC coupling circuit at the receiver front end. This coupling circuit uses 1000 pF capacitor and 150 ohm termination resistance. If a cable carrying a diplexed signal is plugged into another vendor's equipment, the residual signal coupling, described further below, may be quite severe. The critical factor is the value of the coupling capacitance, as larger values increase the unwanted signal coupling. As there is no standard for this circuit, a preferred embodiment will allow for blocking the diplex signal when connecting to an external host device that cannot safely support the present invention. Preferably, a media converter will be designed to overcome such connection problems. Note, however, that this only applies to JBOD configurations ("just a bunch of disks," i.e. no controller for the disks within the JBOD unit) with DB-9 twinax cables.

Regarding equalizer compatibility, in an embodiment of the invention, the low frequency signal is coupled onto the twinax cable with 1k ohm resistors instead of inductors. This places a high DC resistance in series with the RS-232 type low frequency driver. An equalizer, which would be used for a long twinax cable, will have a fairly low impedance between the differential signals. There will be a voltage divider effect between the equalizer impedance and the coupling resistors which will severely attenuate the low frequency signal. This will be a problem if an equalized cable assembly is used. A preferred solution to this issue is to use equalizers for the DB-9 twinax cable, and tapping off the low frequency signal before the equalizer. This would allow the DB-9 to DB-9 media converter to contain an equalizer if we want to support twinax links over 10 meters long. Preferred embodiments will not operate with twinax cables which have an equalizer in the cable connector housing.

Regarding crosstalk, in an embodiment of the invention, the invention will utilize twinax having four conductors contained within a common shield (hereinafter twinax4), wherein isolation between the two conductor pairs is maintained by locating the two pairs orthogonally in the cable. This insures symmetrical coupling from the + and - differential signals, and the net differential crosstalk between pairs is 0. One pair carries the outgoing signals and the other pin carries the returning signals. For differential coupling of the RS-232 type low frequency signal, there will not be any substantial crosstalk coupling from

one pair to the other.

A circuit for the second method, common mode coupling, is shown in FIG. 8. Preferred embodiments of the invention utilize the fact that the Fibre Channel receiver has common mode rejection to reduce noise coupling from the low frequency signals.

- 5 Consequently, since the two signaling approaches on the cable are orthogonal, allowing each receiver to reject the signal from the other. The high frequency signal will preferably be driven differentially on the twinax pair, and the low frequency signal will preferably be driven on both signal conductors with the cable shield serving as the signal return path. The advantages of this method are that it allows potentially higher signaling rates, is
- 10 compatible with equalized cable assemblies, and is potentially more compatible with host interfaces. The noise rejection and crosstalk performance are worse however.

- Regarding ground noise issues, in an embodiment of the invention, any ground potential difference between the transmitter and receiver ends of the cables will couple directly into the common mode low frequency signal as noise. While this is not a new
- 15 problem for RS-232 type links, it is often overcome by running very large signal swings. A preferred embodiment of the invention would be to keep the swing at 5V peak to peak, or less, which may not be as immune to this noise as traditional links.

- Regarding magnetic susceptibility, in an embodiment of the invention, by running the low frequency signal as a common mode signal, the twinax is being used essentially as
- 20 a coaxial cable. This type of cable configuration allows the shield current induced by an AC magnetic field to impose a noise voltage onto the low frequency signal. This voltage will detract from the signal's noise margin.

- Regarding coupling noise, in an embodiment of the invention, coupling of the RS-232 type low frequency signal into the Fibre Channel receiver is now less of a
- 25 problem because the residual signal coupling is being applied as a common mode noise source. The receiver has very good rejection to low frequency common mode noise, so the noise does not create any jitter on the Fibre Channel signal. This allows the low frequency signal pulse filtering to be relaxed, and higher data rates can be accommodated.

- The Fibre Channel receiver has a limited common mode rejection range, so the
- 30 common mode signal must be attenuated before it reaches the receiver. This attenuation is achieved by a voltage divider effect between the Fibre Channel coupling capacitors, and



the 150 ohm termination impedance. Preferred embodiments will have termination mode up of two 75 ohm resistors to VBB or some other DC source to provide a common mode termination.

Regarding external compatibility, in an embodiment of the invention, the  
5 attenuation of the common mode low frequency signal in the Fibre Channel receiver is accomplished by the AC coupling circuit and termination at the receiver front end. The value of the coupling capacitor is less critical in this configuration, as any additional noise will be rejected. The critical factor for interoperability is that the termination be common mode. As there is no standard for this circuit, a preferred embodiment will allow for  
10 blocking the duplex signal when connecting to an external host device that cannot safely support the present invention. Preferably, a media converter will be designed to overcome such connection problems. Note, however, that this only applies to JBOD configurations with DB-9 twinax cables.

Regarding equalizer compatibility, in an embodiment of the invention, the  
15 common mode transmission of the low frequency signal is fully compatible with equalized cable assemblies. The equalizer will act as a low impedance in series with the signal path which will have a minimal effect.

Regarding crosstalk, the common mode low frequency signal does have a crosstalk  
20 issue on the twinax4. The cable has capacitive coupling between the two signal pairs which induces a noise voltage on the other low frequency signal in the cable. This coupling is dependent on risetime, and is sensitive to cable length. In preferred embodiments of the invention, overall the crosstalk will only be a few percent, and will not significantly degrade the noise margin of the signal.

FIG. 9 is a simplified diagram of an embodiment of the duplexing of the invention,  
25 wherein first and second signals are transmitted on a communication link, and wherein the communication link is preferably twin-ax or twisted pair cable having a transmission-end and a receiving-end. The communication link has a pair of conductors including a first conductive path and a second conductive path. The first signal to be transmitted is designated as signal (A), and the second signal is designated as signal (B). Signal (A) is  
30 sent as a differential mode signal. Signal (B) is sent as a common mode signal. The first conductive path contains the addition of signals (A) + (B). The second conductive path

contains the subtraction of  $(B) - (A)$ . At the receiving-end, the first and second conductive paths are both sent separately to an adder and a subtractor. At the subtractor, the second conductive path's signal  $(B)-(A)$  is subtracted from the first conductive path's signal  $(B)+(A)$ , resulting in the  $(A)$  signal. At the adder, the second conductive path's signal  $(B)-(A)$  is added to the first conductive path's signal  $(B)+(A)$ , resulting in the  $(B)$  signal. A useful consequence of this embodiment of the invention is that in the subtractor, common mode line noise is efficiently eliminated, and in the adder, differential mode line noise is also eliminated.

FIG. 10 is a particular implementation of the FIG. 9 embodiment, using transformers.

Referring now to FIGS. 11-13, the common mode coupling embodiment of a diplexer 50 is shown in greater detail. The RS-232 type low frequency signal output is provided to an inverter with hysteresis input 120. The signal is then passed through a low pass filter 122. The low pass filtered signal is then common mode coupled onto the twinax cable. The return signals are summed at a node 124 where both returning signals are connected. The summed signal is low pass filtered by capacitors 126 connected to ground. The summed and filtered signal is provided to a receiver 128 which in the presently preferred embodiment is an LM 393 dual comparator.

Referring to FIG. 13, the expansion connector 52 is shown. A DB9 connector is the presently preferred cable connector 52. Signals received on the connector are passed through an equalizer 130. The received signals are then passed to the low frequency input portion of FIG. 11 and the high frequency input portion of FIG. 12 which will now be described. The high frequency input portion passes the received signals through a high pass filter 132. A port bypass circuit 134 then acts as a differential receiver. The differential receiver acts as a subtractor looking at the difference between the received signals. The resulting high frequency signals are passed through the port bypass circuit 134 to the shunt 40 and to a signal detector 142. The output of the signal detector 142 may be used to allow the shunt 40 to remain in its normal state when no signals are detected. When a signal is detected the shunt 40 may assume its switched state if authorized by the environmental monitor.

The received signals that are passed through the shunt 40 continue onto the serial

bus 24 represented by the FC-SHUNT signals. Outgoing high frequency fibre channel signals are provided by serial bus 22 represented by the FC-LOOP signals. The outgoing signals are driven through the shunt 40. Depending upon the state of the shunt 40, these signals will either loop back to the serial bus 24 on the FC-SHUNT lines or proceed  
5 towards the cable connector 52. The signals sent to the cable connector 52 are differentially driven by shunt 40. One signal is inverted and the other noninverted. The inverted and noninverted signals are passed through a high pass filter 138. Thereafter, the fibre channel signals are combined with the low frequency signals from the low frequency output portion of the diplexer to produce a sum of the high and low frequency signals and  
10 a difference between the low and high frequency signals. The sum and difference signals may then be passed through an equalizer 140 before they reach the cable connector 52.

## WE CLAIM:

1. A dynamically upgradeable disk array chassis comprising:  
a chassis;  
a plurality of ports within said chassis that make connection with a plurality of  
5 data storage devices;  
serial bus means for interconnecting said plurality of data storage devices, said bus  
means including a first bus for passing data in one direction and a second bus for passing  
data in a second direction;  
a shunt connected to the first bus and the second bus and having a normal state in  
10 which the first bus is connected to the second bus across said shunt and a switched state in  
which each of the first bus and the second bus is connected to a separate output;  
an environmental monitor connected to said shunt to control the state of said  
shunt;  
a communication path connected to said monitor;  
15 connection means, on the chassis coupled to said communication path and to said  
shunt, for making data communication connection with a communication path and a serial  
bus from a new disk array chassis, wherein said environmental monitor communicates  
through said communication path with the new disk array chassis before switching said  
shunt to the switched state to connect said serial bus means with the serial bus on the new  
20 disk array chassis.
2. The disk array chassis of claim 1 further comprising a second serial bus  
means, a second environmental monitor, a second shunt, a second communication path  
and a second connection means so as to provide a redundant communication path with the  
25 plurality of storage devices.
3. The disk array chassis of claim 1 wherein said serial bus means carries high  
frequency data signals and said communication path carries low frequency data signals.
- 30 4. The disk array chassis of claim 3 further comprising a diplexer, connected  
between said connection means and both said shunt and said communication path, for

adding the high frequency data signal to the low frequency data signal to produce a sum signal and for subtracting the high frequency data signal from the low frequency data signal to produce a difference signal.

5           5.       The disk array chassis of claim 3 further comprising a diplexer connected to said connection means for receiving a sum signal and a difference signal from the new disk array chassis and including an adder and a subtractor for extracting a high frequency data signal and a low frequency data signal from the new disk array chassis.

10           6.       The disk array chassis of claim 3 wherein said connection means comprises a twinax cable connector.

              7.       The disk array chassis of claim 6 further comprising a diplexer connected to said shunt and said communication path to provide a combination of the high frequency  
15 data signals to said twinax cable connector.

              8.       A dynamically upgradeable disk array chassis comprising:  
a chassis;

              a plurality of ports within said chassis that make connection with a plurality of  
20 data storage devices;

              serial bus means, interconnecting said plurality of data storage devices, for  
carrying high frequency data signals;

              an environmental monitor connected to said plurality of data storage devices to  
detect status of a plurality of conditions within said disk array chassis;

25           a communication path connected to said monitor for carrying low frequency data  
signals;

              a connector accessible outside said chassis to allow communication therethrough  
with a low frequency communication path and a high frequency serial bus from a new  
disk array chassis;

30           a shunt, connected to said serial bus means and having a normal state in which  
said serial bus means is disconnected from said connector and a switched state in which

said serial bus means is in communication, through said connector, with the high frequency serial bus from the new disk array chassis; and

wherein said environmental monitor communicates by low frequency signals through said communication path with the new disk array chassis and afterwards causes  
5 said shunt to switch to the switched state.

9. The disk array chassis of claim 8 further comprising a diplexer connected to said communication path and to said shunt for combining high frequency data signals and low frequency data signals in one direction and for separating high frequency data  
10 signals from low frequency data signals in a second direction.

10. The disk array chassis of claim 9 wherein said diplexer combines by adding the high frequency data signals to the low frequency data signals to produce a sum signal and by subtracting the high frequency data signals from the low frequency data  
15 signals to produce a difference signal.

11. The disk array chassis of claim 9 wherein said diplexer includes an adder and a subtractor for separating the high frequency data signals from the low frequency data signals in the second direction.  
20

12. The disk array chassis of claim 8 further comprising a second serial bus means, a second environmental monitor, a second shunt, a second communication path and a second connector so as to provide a redundant communication path with the plurality of storage devices.  
25

13. The disk array chassis of claim 8 wherein said connector is a twinax cable connector and wherein the high frequency data signals and the low frequency data signals travel through said connector when said shunt is in the switched state.

30 14. A dynamically upgradeable disk array chassis comprising:  
a chassis;

-21-

a plurality of ports within said chassis that make connection with a plurality of data storage devices;

serial bus means for interconnecting said plurality of data storage devices, said serial bus means including a first bus for carrying high frequency data signals in one direction and a second bus for carrying the high frequency data signals in a second direction;

a shunt connected to the first bus and the second bus and having a normal bypass state in which the first bus is connected to the second bus across said shunt and a switched state in which each of the first bus and the second bus is connected to a separate output;

10 an environmental monitor connected to said shunt to control the state of said shunt;

a communication path connected to said monitor for carrying low frequency data signals;

a diplexer connected to said communication path and to the separate outputs of said shunt for combining the high frequency data signals and the low frequency data signals in the one direction and for separating a high frequency data signal from a low frequency data signal in the second direction; and

a connector, connected to said diplexer, accessible outside said chassis to make data communication connection with a low frequency communication path and a high frequency serial bus from a new disk array chassis, wherein said environmental monitor communicates with low frequency signals through said communication path with the new disk array chassis before switching said shunt to the switched state to connect said serial bus means to the high frequency serial bus on the new disk array chassis.

25 15. The disk array chassis of claim 14 wherein said diplexer combines by adding the high frequency data signals to the low frequency data signals to produce a sum signal and by subtracting the high frequency data signals from the low frequency data signals to produce a difference signal.

30

-22-

16. The disk array chassis of claim 14 wherein said diplexer includes an adder and a subtractor for separating the high frequency data signals from the low frequency data signals in the second direction.

5 17. The disk array chassis of claim 14 further comprising a second serial bus means, a second environmental monitor, a second shunt, a second communication path, a second diplexer and a second connector so as to provide a redundant communication path with the plurality of storage devices.

10 18. The disk array chassis of claim 14 wherein said connector is a twinax cable connector and wherein the high frequency data signals and the low frequency data signals travel through said connector when said shunt is in the switched state.

15 19. A method for dynamically upgrading a data storage system comprising:  
providing a disk array chassis having a plurality of data storage devices interconnected by a high frequency data bus, an environmental monitor connected to a low frequency communication serial and a diplexer for combining high frequency signals from the high frequency data path with low frequency signals from the low frequency communication path;  
20 connecting a cable between the diplexer of said disk array chassis and a connector on said data storage system;  
communicating with said disk array chassis with low frequency signals through said cable and said communication bus; and  
switching a shunt in the data storage system to permit high frequency signals to  
25 communicate between said data storage system and said disk array chassis.

49122






FIG. 1

2/9

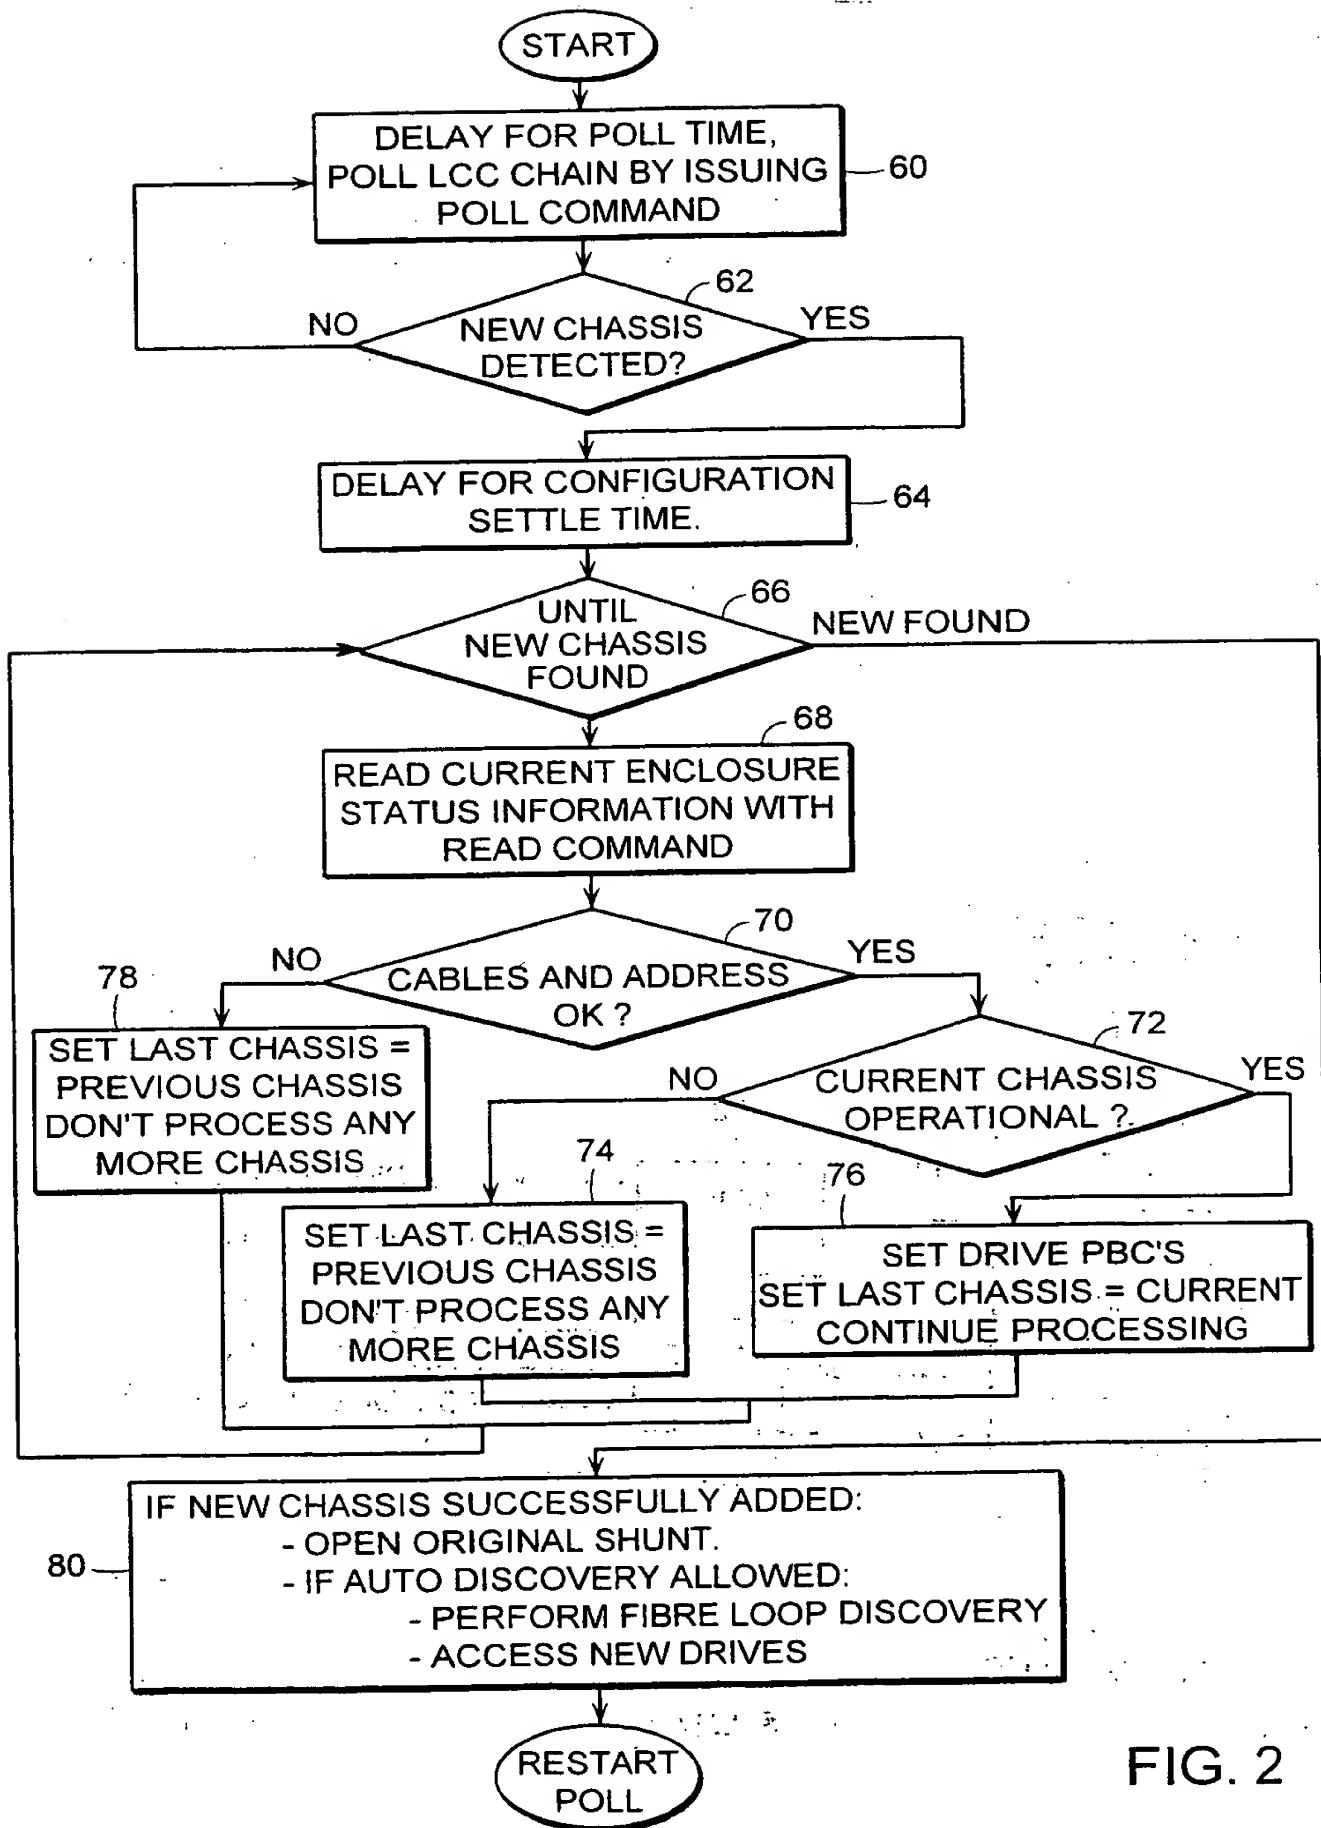


FIG. 2

3/9




FIG. 3




FIG. 4

4/9




FIG. 5




FIG. 6

5/9

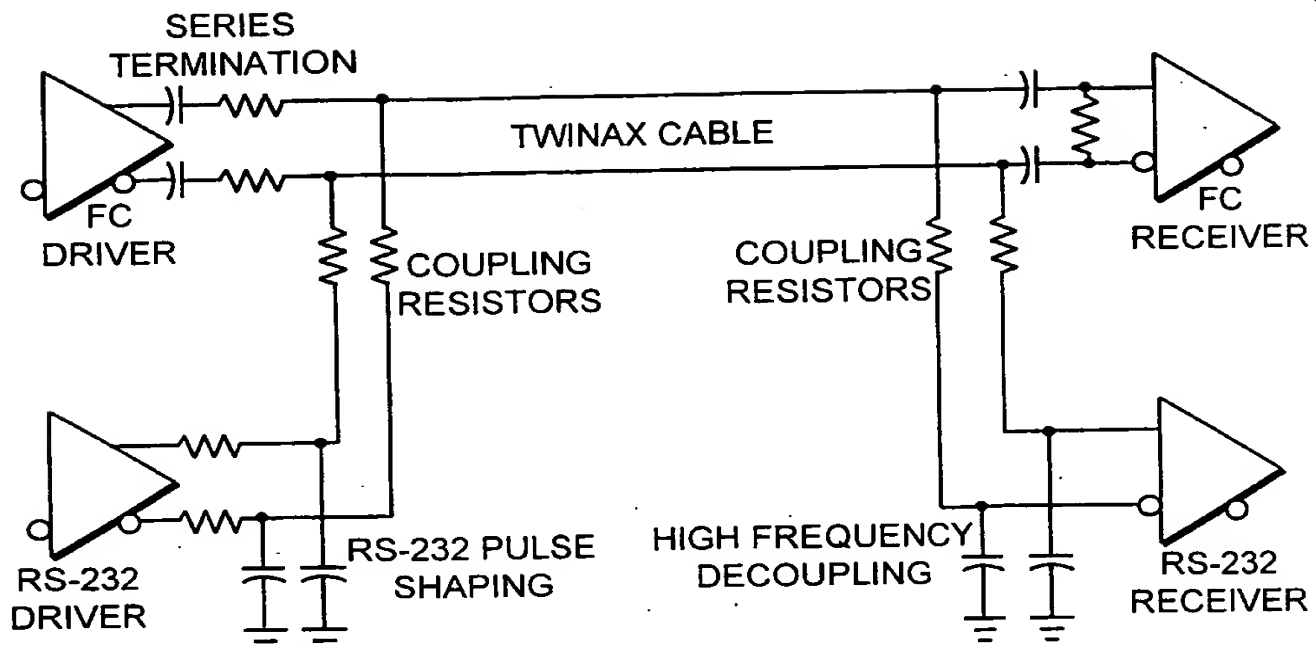


FIG. 7

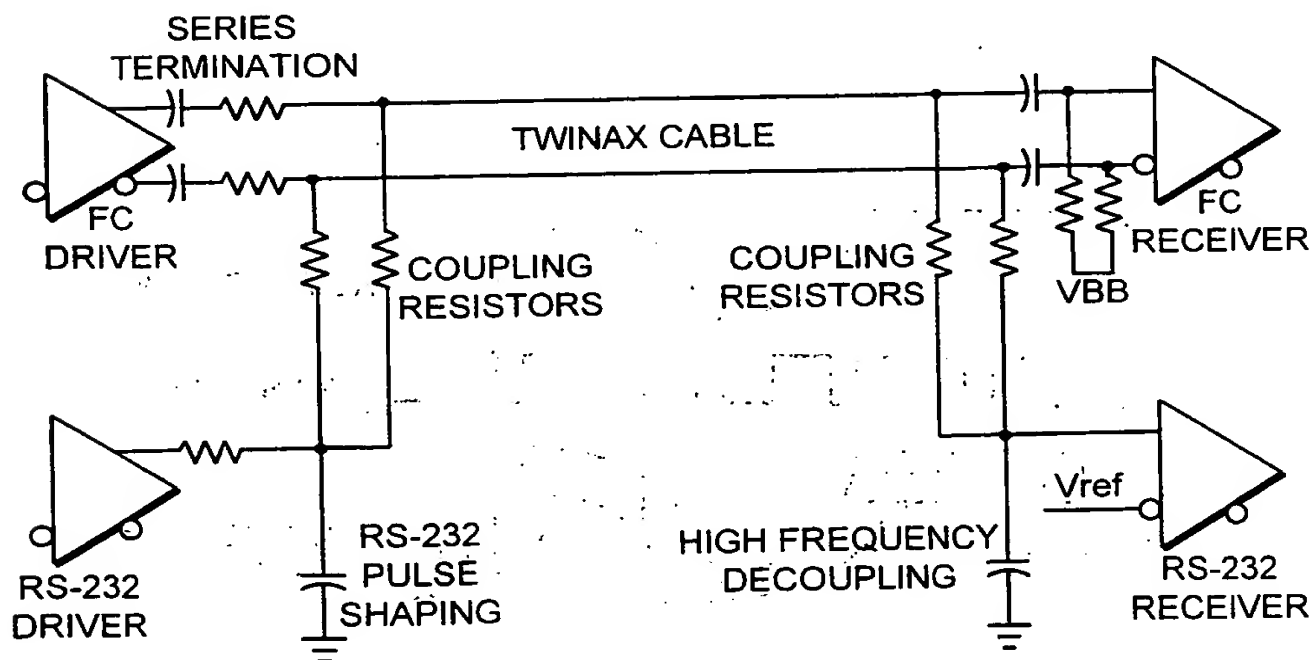


FIG. 8

6/9




FIG. 9




FIG. 10

7/9





FIG. 11

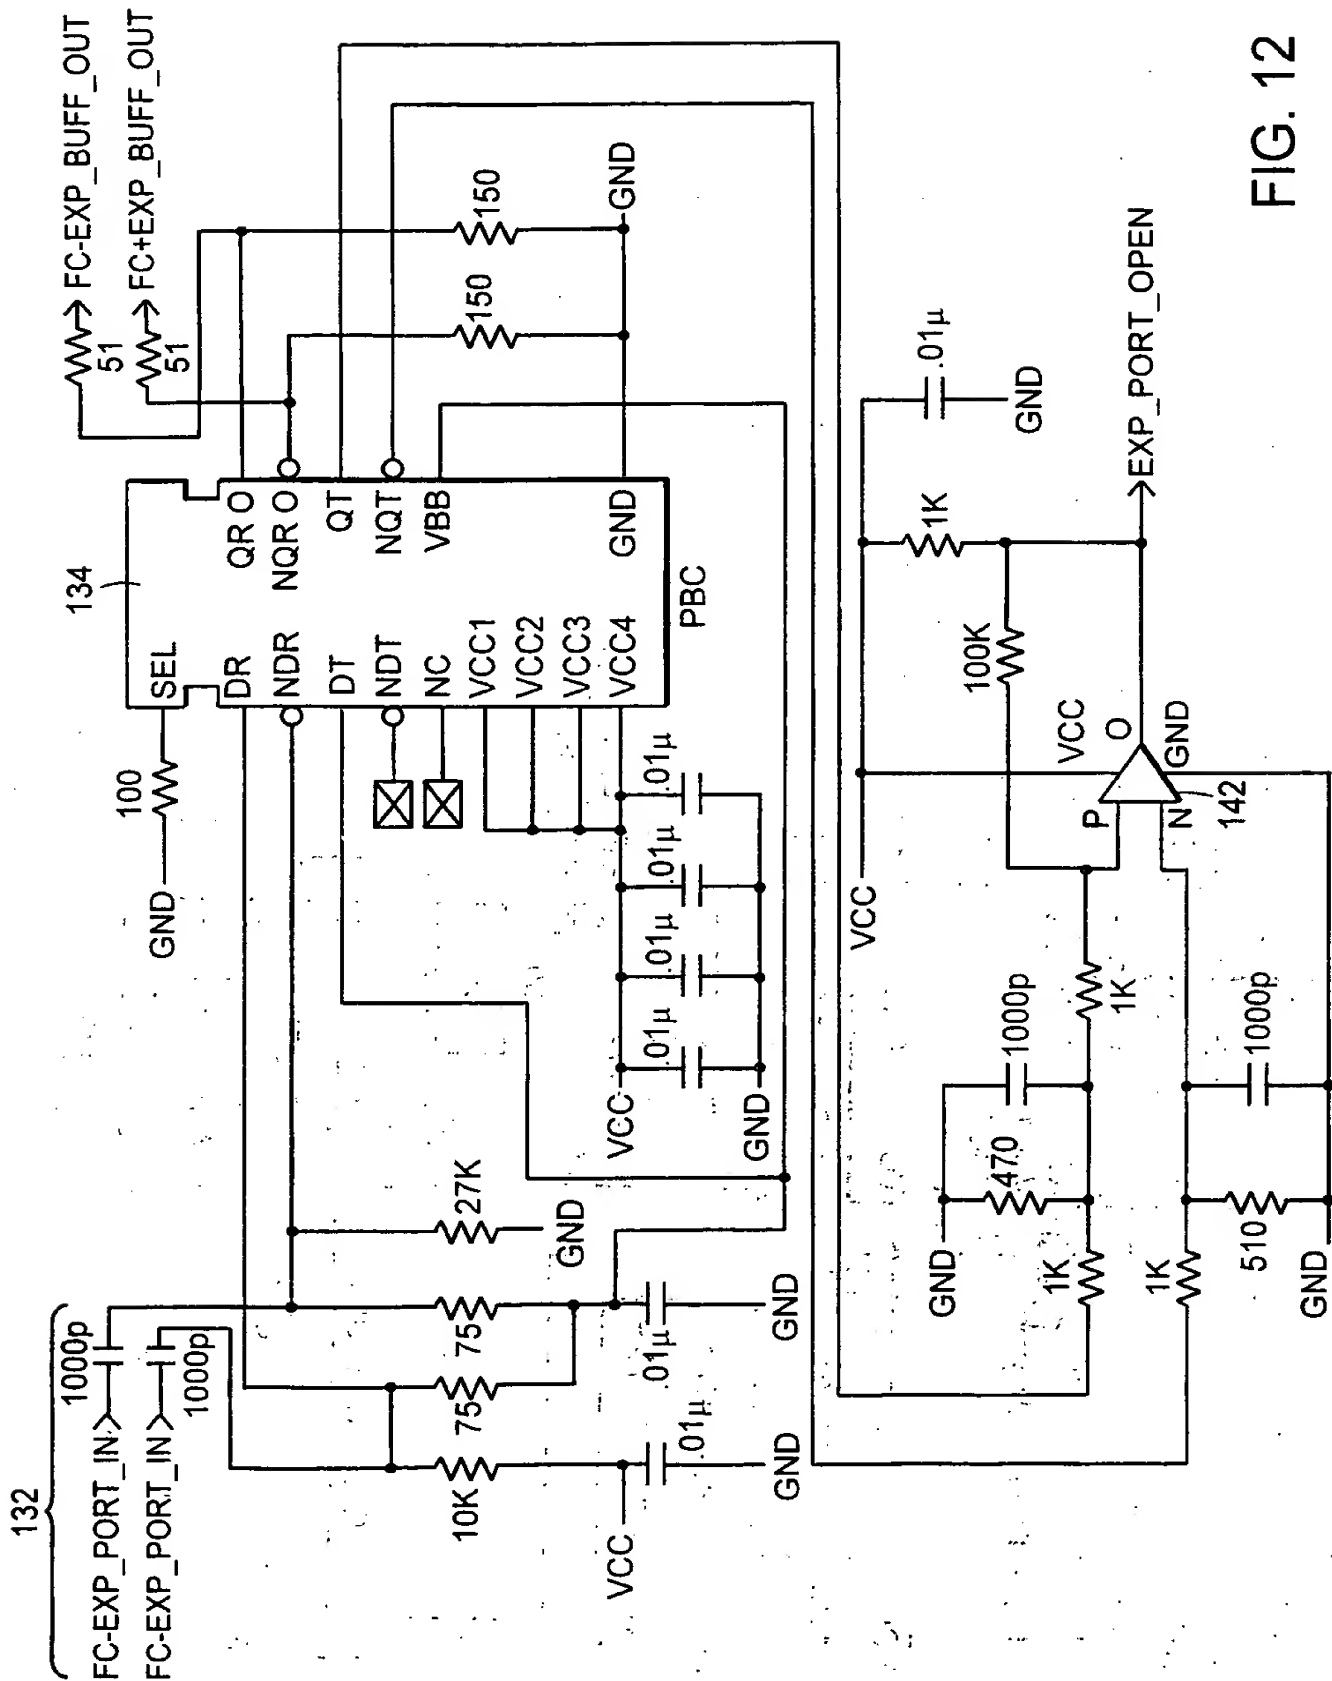


FIG. 12



9/9

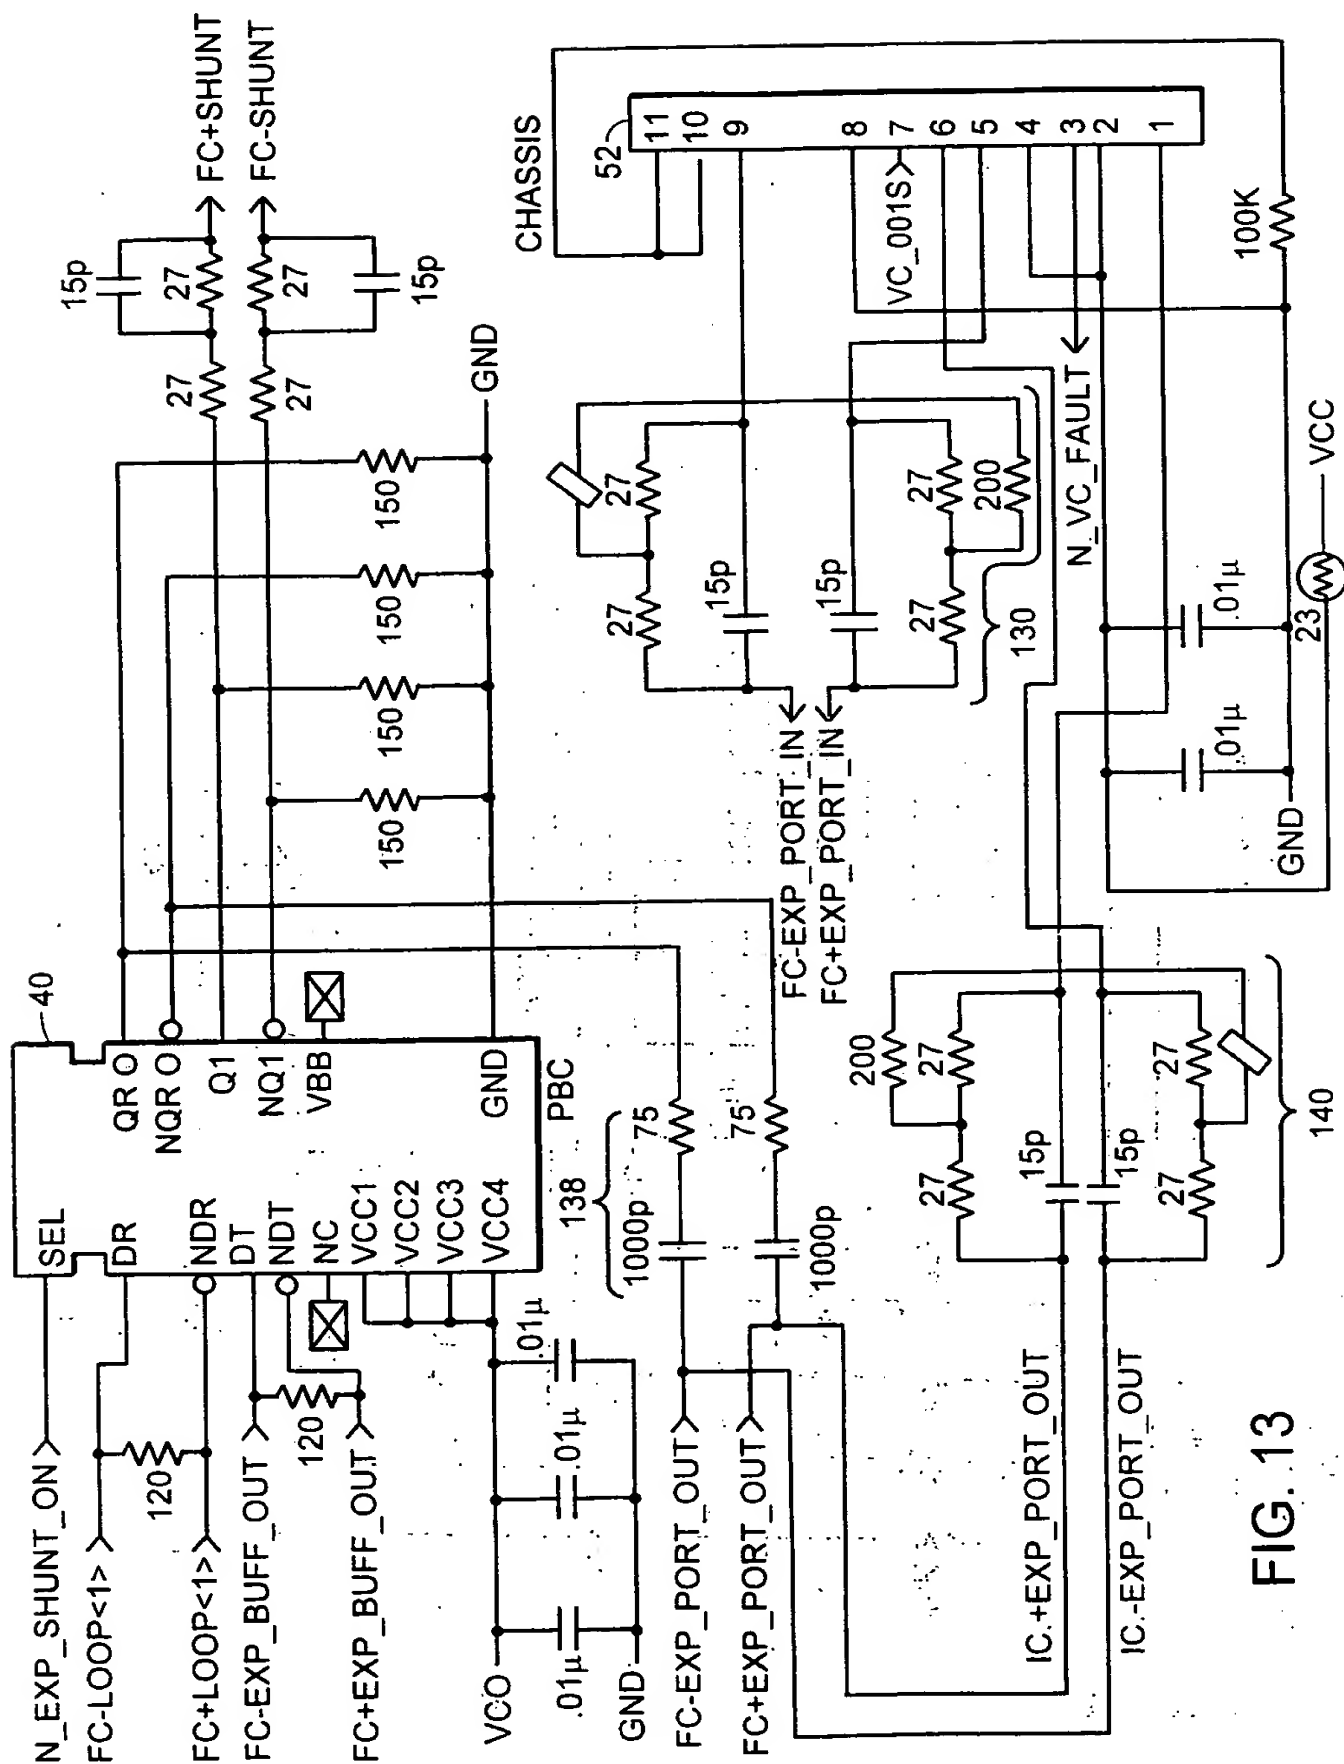


FIG. 13

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US97/17432

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 13/00, 11/20

US CL :711/114, 170; 395/182.05, 283, 311

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 711/2, 4, 5, 112, 114, 115, 170; 395/182.03, 182.04, 182.05, 283, 311, 653

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
Please See Extra Sheet.

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,390,187 A (STALLMO ET AL) 14 February 1995 (14.02.95), see entire document.	1-19
A	US 5,564,116 A (ARAI ET AL) 08 October 1996 (08.10.96), see entire document.	1-19
A	WO 95/13581 A2 (CONNER PERIPHERALS, INC.) 18 May 1995 (18.05.95), see entire document.	1-19
A	SUTTON ET AL, Fibre Channel: The Digital Highway Made Practical. Seagate Technology Paper. 22 October 1994, pages 1-6.	1-19



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
*A* document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
*E* earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
*L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*G* document member of the same patent family
*O* document referring to an oral disclosure, use, exhibition or other means	
*P* document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

30 JANUARY 1998

Date of mailing of the international search report

06 MAR 1998

Name and mailing address of the ISA/US  
Commissioner of Patents and Trademarks  
Box PCT  
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

GLENN GOSSAGE

Telephone No. (703) 305-3820

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US97/17432

## B. FIELDS SEARCHED

Electronic data bases consulted (Name of data base and where practicable terms used):

USPTO Automated Patent System (APS), files USPAT, JPOABS, EPOABS

search terms: disk, array, chassis, upgrade, dynamic, fibre, channel, shunt, switch, serial, monitor

**THIS PAGE BLANK (USPTO)**